



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ



I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

Dated 29 August 2001



The Patent Office


1/77

Patents Act 1977

Rule 16

23JUL01 E647244-1 D00611
P01/7700 0.00-0117876.3

Request for grant of a patent



The Patent Office
Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

1.	Your reference	GB920010032GB1		
2.	Patent application number (The Patent Office will fill in this part)	20 JUL 2001	0117876.3	
3.	Full name, address and postcode of the or of each applicant (<i>underline all surnames</i>)	INTERNATIONAL BUSINESS MACHINES CORPORATION Armonk New York 10504 United States of America		
	Patents ADP number (<i>if you know it</i>)	00519637001		
	If the applicant is a corporate body, give the country/state of its incorporation	State of New York United States of America		
4.	Title of the invention	A METHOD, SYSTEM AND COMPUTER PROGRAM FOR CONTROLLING ACCESS IN A DISTRIBUTED DATA PROCESSING SYSTEM		
5.	Name of your agent (<i>if you have one</i>)	S R Davies		
	"Address for Service" in the United Kingdom to which all correspondence should be sent (<i>including the postcode</i>)	IBM United Kingdom Limited Intellectual Property Department Hursley Park Winchester Hampshire SO21 2JN		
	Patents ADP number (<i>if you know it</i>)	06449425001		
6.	If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (<i>if you know it</i>) the or each application number	Country	Priority App No (<i>if you know it</i>)	Date of filing (<i>day/month/year</i>)
7.	If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application	No of earlier application	Date of filing (<i>day/month/year</i>)	

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:
a) any applicant named in part 3 is not an inventor, or
b) there is an inventor who is not named as an applicant, or
c) any named applicant is a corporate body.)
- Yes

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description 17

Claim(s) 4

Abstract 1

Drawing(s) 5 XS

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77) 3

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)


11. I/We request the grant of a patent on the basis of this application

Signature

19 July 2001
Date

S R DAVIES

12. Name and daytime telephone number of person to contact in the United Kingdom
- A Sekar
01962 818169



7/77

Patents Act 1977

Rule 15

Statement of inventorship and of right to grant of a patent

The Patent Office
Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

1.	Your reference	GB920010032GB1
2.	Patent application number (if you know it)	20 JUL 2001 0117876.3
3.	Full name of the or of each applicant	INTERNATIONAL BUSINESS MACHINES CORPORATION
4.	Title of invention	A METHOD, SYSTEM AND COMPUTER PROGRAM FOR CONTROLLING ACCESS IN A DISTRIBUTED DATA PROCESSING SYSTEM
5.	State how the applicant(s) derived the right from the inventor(s) to be granted a patent	By employment and agreement
6.	How many, if any, additional Patents Forms 7/77 are attached to this form?	
7.	I/We believe that the person(s) named over the page (and on any extra copies of this form) is/are the inventor(s) of the invention which the above patent application relates to.	
	Signature	19 July 2001 Date
	S R DAVIES	
8.	Name and daytime telephone number of person to contact in the United Kingdom	A Sekar Tel: 01962 818169

Patents Form 7/77

Enter the full names, addresses and postcodes of the inventors in the boxes and underline the surnames

HUGHES, Jeremy Peter James
UK resident
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (if known)

08193856001

TATE, Richard Peter James
UK resident
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (if known)

08193864001

If there are more than three inventors, please write their names and addresses on the back of another Patents Form 7/77 and attach it to this form

REMINDER

Have you signed the form?

Patents ADP number (if known)

A METHOD, SYSTEM AND COMPUTER PROGRAM FOR CONTROLLING ACCESS IN A
DISTRIBUTED DATA PROCESSING SYSTEM

5

FIELD OF THE INVENTION

The present invention relates to access control in a distributed data processing system.

10

BACKGROUND OF THE INVENTION

In the World Wide Web (WWW) environment, client machines communicate with Web servers using the Hypertext Transfer Protocol (HTTP). The web servers provide users with access to files such as text, graphics, images, sound, video, etc., using a standard page description language known as Hypertext Markup Language (HTML). HTML provides basic document formatting and allows a developer to specify connections known as hyperlinks to other servers and files. In the Internet paradigm, a network path to a server is identified by a Uniform Resource Locator (URL) having a special syntax for defining a network connection. So called web browsers, for example, Netscape Navigator (Netscape Navigator is a registered trademark of Netscape Communications Corporation) or Microsoft Internet Explorer, which are applications running on a client machine, enable users to access information by specification of a link via the URL and to navigate between different HTML (web) pages.

When the user of the web browser selects a link, the client issues a request to a naming service to map a hostname (in the URL) to a particular network IP (Internet Protocol) address at which the server is located. The naming service returns an IP address that can respond to the request. Using the IP address, the web browser establishes a connection to a server. If the server is available, it returns a web page. To facilitate further navigation within a web site, a web page typically includes one or more hypertext references known as "anchors" or "links".

Today, there exists a vast amount of web pages whereby information within the web pages is dynamic, decentralised and diverse. For a user, the task of traversing the information can be very difficult and time-consuming. Therefore, there is a need for an efficient and automated method of traversing this information, so that a user is able to find relevant information amongst the vast amount of pages that exist.

A "robot" is a type of "agent" that is one solution to this problem. An agent is a computer program that is goal-oriented, that is, an agent tries to achieve some end result. For example, an agent could perform a task on behalf of a user and this is shown in FIG 1, by using the example of the Internet. In FIG 1, a user at a client computer (100) dispatches two agents via a controlling application program running on the client (100). "Agent 1" and "Agent 2" are dispatched over a network (110), which in this example, is the Internet. Since agents can be customised, the user can dispatch "Agent 1" to find a first piece of information held on a remote server (120), for example, the address of the nearest pizza restaurant. The user can also dispatch "Agent 2" to find a second piece of information, for example, the phone number of a taxi firm, which in this example is also held on the same remote server (120).

A robot is a special automated form of agent. The robot may simply react to changes in its environment, or when subjected to stimuli. "Web" robots are widely used for search and extraction of information held in web pages. They also have other uses, such as for personal shopping, whereby the robot collects information about products and prices from the WWW and presents this to the user. Robots can also be utilised in other mediums, such as, in databases.

Information gathering robots, typically used to retrieve unstructured information, such as text or images, are also known as "spiders", "crawlers" or "wanderers". These types of robots are most often used in highly interconnected data environments, such as the WWW. The term "crawling" is often used to denote the process of moving through an environment in a managed way. Specifically, an information gathering robot is a program that automatically explores the WWW by retrieving a document and recursively retrieving some or all of the documents that are linked to it. The robot has thus generated a web index of documents.

There are two main categories of crawling, namely, unfocussed and focussed. In unfocussed crawling, the robot is not looking for anything in particular and its main aim is to gather as much information as possible. This technique is often used by a "search engine", which searches through a web index in order to help locate information by keyword for example. Focussed crawling indicates that the robot is looking for a particular piece of information. This technique is used by a specialised robot such as a shopping robot.

More information about agents and web robots can be found in the book "Internet Agents: Spiders, Wanderers, Brokers and Bots" by Fah-Chun Cheong, New Riders Publishing, 1996.

Many robots are used for legitimate reasons, such as, for searching. Robots are often developed by well-known organisations, for example, search engine technology from Yahoo, Lycos, Google and so forth. However, when the first robots were developed, they had a reputation for sending hundreds or thousands of requests to each web site when gathering documents and this often resulted in the web site being overloaded. Although the development of robots has improved, some robots may still exhibit unfriendly behaviour, and it is this type of behaviour that an administrator may not be willing to tolerate.

Another reason for an administrator to want to block access to robots is to prevent them from indexing dynamic information. Using the example of searching again, many search engines will use information collected from a web site repeatedly, for weeks or months to come. Obviously, this feature is not much use if the web site is providing stock quotes, news, weather reports or any other information that will be out of date by the time a user finds it via a search engine. Other malicious robots are routinely used to systematically copy content assets from public web sites.

Currently, there are a number of methods of excluding robots from web sites. One example is the "Standard for Robot Exclusion" proposed by Martijn Koster and available at <http://www.robotstxt.org/wc/wxclusion-admin.html>. The protocol specifies a format for a file "Robots.txt", located in a web server's root directory. This file provides a means to request that a named robot limits its activities at a particular web site, or requests that a robot leave a web site. In FIG 2, the first line in the robots.txt file (200) identifies that the exclusion policies refer to a robot called "Robot_1". The second line of the file (200) specifies that Robot_1 should not visit any URLs where "/england/london" is present after the host name in the URL, where a host name may take the form "www.corp.com". In the third line, the robot is also excluded from visiting any URLs where "/france/paris" is present after a host name.

However, the disadvantage with the Standard is that the exclusion policies may or may not be obeyed. This is because, although a robot may review the robots.txt file, it is the decision of the robot's creator as to whether or not the file is obeyed. In the case of malicious robots, the Standard is often ignored or misinterpreted, resulting in web sites being adversely affected by the actions of uncontrolled robots. If this occurs, a major challenge for administrators is to identify malicious robots and put in place manual methods for explicitly dealing with them promptly and effectively.

Some robots may be relatively simple to detect, since their activity may be concentrated into a short time period. Alternatively the robot may manifest itself as a form of "denial-of-service" or "ping attack". In this case a server is repeatedly hit by requests therefore limiting its capability to respond effectively. However, other robots use techniques so that they cannot be detected easily. One example is by hiding amongst the "noise" of traffic created by legitimate users of the system. Another example is by taking hours to complete a navigation of a system. In these cases, the manual and explicit exclusion of robots is difficult and unreliable.

Another method of controlling robots, or spiders in the case of this method, can be found at <http://www.spiderhunter.com>. The method described at this web site uses data collected when a user visits a web site, rather than using analysis of log files. To collect data, the method utilises three pieces of information, namely, an IP address associated with the user, the name of the spider being used and the file being requested. The method uses a neural net to check for new information and compares the new information against known information. For example, an IP address of a potential spider is checked to see whether it matches a known IP address of a spider. The neural net uses a baseline to determine whether the user is legitimate and uses weights to determine the likelihood of the user being a spider.

There are many disadvantages with using a neural net for detection of robots. For example, the output results from this method will only be as accurate as the amount of information input into it. Also, an administrator will not be able to modify the underlying detection method to suit their needs, rather, only the weights can be modified. This particular method also relies on the fact that a potential spider provides an IP address, however, if a spider enters a site through multiple proxies, it may be able to hide its IP address. Another web site offering a similar service is "Spider Central" which can be found at <http://www.john.php4hosting.com>.

Therefore there is a need for a method of automatically detecting and managing malicious robots, so that administrators can control access to their web sites, servers and systems more effectively.

DISCLOSURE OF THE INVENTION

According to a first aspect, the present invention provides a method of controlling access to information in a distributed data processing system having: a server for storing said information, in which said server further comprises a logging tool for creating a log file, and a client

computer comprising an application program for controlling a software agent, in which said software agent requests said information from said server, said method comprising the steps of: identifying a software agent; in response to said identifying step, storing all requests from said identified software agent in said log file; in response to said storing step, analysing said log file; in response to said analysing step, monitoring behaviour of said identified software agent, and in response to said monitoring step, invoking at least one of a plurality of pre-defined rules to control said behaviour of said identified software agent.

Preferably, the information is represented within any number of a plurality of web pages, whereby each web page has a non-visible hypertext link inserted into their main bodies. Preferably, when a software agent requests one of a plurality of web pages, the agent is identified. The identifying step further comprises the steps of: dynamically generating a first unique identifier; dynamically inserting the first unique identifier into a non-visible link associated with one of a plurality of web pages, and determining whether the one of a plurality of web pages is associated with further of a plurality of web pages. In the preferred embodiment of the present invention, the software agent is a Web robot.

Preferably, if it is determined that the one web page is associated with further web pages, the first identifier is dynamically inserted into further non-visible links. The further non-visible links are of course associated with the further web pages. In the preferred embodiment of the present invention, the first identifier is a session ID, which will be described in more detail later on.

Preferably, if it is determined that the one web page is not associated with further web pages, the one page is firstly sent to the identified software agent. Then, the identified software agent requests any number of a plurality of links associated with the one web page from the server. In response to the request, the identified software agent extracts the any number of a plurality of links. Next, the identified software agent passes the any number of a plurality of links to the client application program. Once the client application program has received the links, it decides which of the links to display, based on its own logic. More preferably, the any number of links are displayed within a web browsing session running on the client computer.

Preferably, the analysing step of the method described above, further comprises the step of: identifying a first value associated with the any number of a plurality of web pages and a second value associated with the further of a plurality of web pages. Preferably, the monitoring step of the

method described above, further comprises the steps of: utilising the first and second values to generate a third value, in which the third value is associated with the identified software agent, and utilising said third value and a fourth value associated with said all requests, to associate said first identifier and said identified software agent with a fifth value, in which said fifth value is associated with a probability.

More preferably, if the third value is not more than or equal to a first pre-determined threshold, the log file is analysed further. More preferably, if the fourth value is not more than or equal to a second pre-determined threshold, the log file is analysed further. In the preferred embodiment, the thresholds are pre-determined by an administrator.

Preferably, prior to the invoking step, the identified software agent is associated with a profile, whereby the profile comprises any number of a plurality of data fields unique to the identified software agent. Preferably, the invoking step further comprises the steps of: associating the any number of a plurality of data fields with the at least one of a plurality of pre-defined rules; determining whether a pre-defined response needs to be activated, and in response to a successful determining step, activating the pre-defined response. Advantageously, the administrator can configure the pre-defined responses to suit the environment.

Preferably, the log file further stores an address associated with a software agent and a name associated with the software agent. More preferably, at least one of the data fields is extracted from the log file. This is advantageous, in that the information captured about the software agent can now be re-used.

Preferably, a second unique identifier is also generated and the inserted into any non-visible link. In the preferred embodiment of the present invention, the first identifier is a page ID, which will be described in more detail later on. Preferably, the distributed data processing system further comprises an application server. In the preferred embodiment, the application server processes the techniques.

Preferably, the at least one of a plurality of pre-defined rules control a plurality of thread priorities associated with the server, in which at least one of a plurality of threads is associated with a software agent. Advantageously, the preferred embodiment ensures that a robot does not consume too much of a system's resources at the expense of other users.

According to a second aspect, the present invention provides a system for controlling access to information, for use in a distributed data processing system, said distributed data processing system comprising: a server for storing said information, in which said server further comprises a logging tool for creating a log file, and a client computer comprising an application program for controlling a software agent, in which said software agent requests said information from said server, said system for controlling access comprising: means for identifying a software agent; means, responsive to said identifying means, for storing all requests from said identified software agent in said log file; means, responsive to said storing means, for analysing said log file; means, responsive to said analysing means, for monitoring behaviour of said identified software agent, and means, responsive to said monitoring means, for invoking at least one of a plurality of pre-defined rules to control said behaviour of said identified software agent.

According to a third aspect, the present invention provides a distributed data processing system comprising: a server for storing said information, in which said server further comprises a logging tool for creating a log file, and a client computer comprising an application program for controlling a software agent, in which said software agent requests said information from said server, and a system for controlling access to information, as described above.

According to a fourth aspect, the present invention provides a computer program product comprising computer program code stored on a computer readable storage medium which, when executed on a computer, performs the method as described above.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will now be described, by way of example only, with reference to preferred embodiments thereof, as illustrated in the following drawings:

FIG 1 shows a prior art overview of a distributed data processing system utilising agents;

FIG 2 shows a prior art representation of a file for exclusion of web robots;

FIG 3 shows a prior art distributed data processing system in which the present invention may be implemented;

FIG 4 is a flow chart showing the operational steps involved in the robot activity logging technique, according to the present invention;

FIG 5 is a flow chart showing the operational steps involved in the robot activity monitoring technique, according to the present invention; and

FIG 6 is a flow chart showing the operational steps involved in the robot management technique, according to the present invention.

DETAILED DESCRIPTION OF THE INVENTION

FIG 3 shows a pictorial representation of a distributed data processing system in which the present invention may be implemented. Distributed data processing system 300 comprises a number of computers, connected by a network 302. Server 304 is preferably a web server and is connected to network 302 along with storage unit 306 and clients 308, 310 and 312. In the depicted example, distributed data processing system 300 is the Internet, with network 302 representing a world-wide collection of networks and gateways that use the transmission control protocol over internet protocol (TCP/IP) suite of protocols to communicate with one another.

Generally, the present invention provides a technique for identifying the existence of a robot and tracking its behaviour on a web site. Once identified, a management process is put into place to either control the robot or remove it entirely. The present invention will be described under three main categories:

1. A robot activity logging technique
2. A robot activity monitoring technique
3. A robot management technique

1. A robot activity logging technique

The aim of this logging technique is to facilitate the identification of robot activity on a web site. It should be understood that the term "robot" encompasses spiders, wanderers etc in this description.

This technique will now be described with reference to FIG 4. The pre-requisite for this logging technique is a "session ID". A session ID is a unique identifier as it has a different value each time a web browser is re-started. Preferably, web pages supporting this technique are generated

5 dynamically and make use of URL-rewriting session ID management. A session ID is generated (step 400) by a session management system located in the server (304). It should be understood that although dynamic session IDs are used in the preferred embodiment, any other form of unique identifier may be used.

10 Preferably, all web pages on a web site supporting this technique have a prior art non-visible hypertext link inserted (step 405) into their main bodies. This is called a "hidden link". A hidden link is a URL that cannot be seen on the web page by a user, that is, there is no hypertext visible. However, a valid hypertext ("HREF") reference is present so that a search engine, for example, can find it.

15 Preferably a session ID is dynamically inserted into a hidden link when a web page is generated. This resulting information is preferably stored in a log file. Preferably, two sets of information are added (step 410) to the hidden link, namely:

- 20 1. A "page ID" - this is a unique identifier for every web page on a web site. In the examples below it is represented by a string of numbers.
- 25 2. A "session ID" - as described above, this identifier is dynamically inserted into the hidden link when the web page is generated. In the examples below it is represented by a string of numbers.

Both sets of information are encoded into the hidden link, so that a web application server or a web server plug-in can extract the information when required.

30 An example of the HREF attribute is provided below:

```
<A HREF="/valid_dir/valid_page0012x1234.html"></A>
```

35 In this example, "0012" is a page ID and "1234" is a session ID. Note that a session ID or page ID can either be added to the hidden link as a standard parameter, for example (using a session ID):

```
http://www.mysite.com/blindlink15?sessionID=1234
```

40 or, the identifiers could be embedded within the page name, for example (using a session ID):

```
http://www.mysite.com/blindlink15_1234.html.
```

Referring to the addition technique, advantages include the fact that individual robots can be tracked. More specifically, the client application program that controls a robot is tracked. Referring to the embedding technique, the advantage is that session IDs will not be lost if parameters are stripped from a URL. However, the web server or application server (which is located below the web server in the system architecture) will need to do more work in order to extract the information in the first place. The advantage of associating session IDs with URLs, rather than using cookies on a client computer to generate a session ID, is that robots often ignore cookies.

Sometimes robots are configured to ignore hidden links that point to a web page with no on-screen text and therefore these robots cannot be logged. The HREF attribute will not contain any visible text and may look like this:

```
<A HREF="/abc.html"></A>
```

Preferably, to overcome this problem one or more characters, which are set to the background colour of that web page, are used between HTML tags. The HREF attribute may look like this:

```
<set foreground to be the same as background><A HREF="/abc.html">  
abc/A><next foreground colour>
```

Therefore, the robot is forced to follow this hidden link and can then be logged.

Preferably, if other links on the web page lead to further dynamic pages, that is, a positive result to step 415, these links are similarly encoded (step 417) with the same session ID, but not with the same page ID. This is because the further dynamic pages are navigated to within the same session. However, each further dynamic page will be associated with a unique page ID. In this description, a web page containing a hidden link is termed a "hidden page" and a web page identified by a hidden link is termed a "blind page".

The hidden link is not visible on a client web browser and therefore it is unlikely that typical users can navigate to a web page that the hidden link refers to by using standard methods. It can therefore be assumed that anything requesting the web page is extracting the hidden link by a different means, such as by page parsing which is typically used by spider systems. In this process, the spider utilises the HTML source for

the web page which a web browser uses to display content, and extracts HTML code that represents links to other web pages.

Referring back to FIG 4, if there are no further links to further web pages, that is, a negative result to step 415, the full web page is downloaded to the robot and it will then extract (step 420) all links on the page, including the hidden links. It will then pass (step 425) all the links back to the controlling client application program. This client application program will then decide (step 430) on which links to follow based on its own logic. For example, if the program specified to the robot that it should only return links relating to shopping, the program may decide to follow only these links.

Preferably, the web server continuously logs (step 435) all HTTP requests made to it and this information is stored (step 440) in a log file. A HTTP request may be structured as shown below:

`http://www.mysite.com/blindlink15?sessionID=1234`

Preferably, the log file also stores the IP address of the requester and the name of the user, that is, a "user_agent" field.

At step 440, the logging technique may either finish if no more processing is to be done, or, at "A", the processing passes to FIG 5.

2. A robot activity monitoring technique

The aim of this monitoring technique is to track the behaviour of a robot on a web site.

This technique will now be described with reference to FIG 5. The pre-requisite for this technique is a web server log file containing data captured by the activity logging technique as described above. The log file is input into the processing at "A".

In this technique, firstly a robot identifier is automatically determined by analysing (step 500) the log file to identify (step 505) all blind pages associated with a session ID. In practice, this may occur by parsing the HTTP requests stored in the log file in order to extract the values relating to the blind pages and session IDs. For example, for a given session ID "1234", the blind pages requested are "0012", "0034", "0045", "0027", "0067", "0094", "0101", "0002", "0032", "0039".

The session ID is then given (step 510) a score, termed a "robot score". The score is calculated by taking the total number (x) of blind pages requested as a percentage of the total number (y) of hidden pages requested.

$$\text{ROBOT_SCORE} = \frac{x \text{ Blind Pages}}{y \text{ Hidden Pages}} * 100$$

Once a robot score has been calculated, it is utilised (step 515) by a "robot identification rule", which is used to establish whether or not a session ID can be associated with a robot with a probability value. The rule takes the form:

```
IF ROBOT_SCORE >= THRESHOLD1
AND NUMBER_OF_REQUESTS >= THRESHOLD2
THEN ASSOCIATE SESSIONID WITH ROBOT WITH PROBABILITY_P
```

In the above rule, the "NUMBER_OF_REQUESTS" parameter represents the total number of requests received during the monitored session. That is, the requests for the hidden pages, blind pages and any other pages requested. "THRESHOLD1" and "THRESHOLD2" are pre-defined values set by an administrator. The purpose of the thresholds is to reduce the probability of false alerts. For example, if it is realised that a human user shows signs of robot activity, such as, if that user makes a large number of requests for hidden pages, an administrator can vary THRESHOLD2 (e.g. by raising the pre-defined value for THRESHOLD2). Therefore, the monitoring process is flexible enough to be implemented in different environments.

The first threshold is based on the hypothesis that the probability of a typical user requesting every blind page when viewing a hidden page is extremely small, whereas a robot system is much more likely to exhibit this behaviour. THRESHOLD1 is therefore set to a high value, such as, "98". If the requirements for the first threshold are met, that is, a positive result to step 520, then processing passes to the second threshold.

The second threshold acts as a sanity check, whereby the meaningfulness of the robot score increases as the number of requests made increases. THRESHOLD2 is therefore set to a minimum value, such as, "10". If the requirements for the second threshold are met, that is, a positive result to step 525, then processing passes step 530.

If the requirements for the two thresholds are not met, processing passes to step 535. In this step, if more information is to be analysed, that is, a positive result to step 535, processing passes once more to step

500. However, if there is no more information to be analysed, that is, a negative result to step 535, processing finishes.

Referring back to step 530, a probability value (P) is simply calculated as shown below:

$$P = \text{ROBOT_SCORE} / 100$$

The probability value represents the likelihood of the monitored behaviour being that of a robot. The range of values is from 0 to 1. That is, a malicious robot will have a higher probability value than a non-malicious robot.

By using a rule-based paradigm it is easy to add extra elements to the robot monitoring process. For example, the time taken between requests for a hidden page and its associated blind page may be included as another feature within the rule. This flexibility allows the rules to be optimised and enhanced by an administrator based on trials and experience, if desired.

The monitoring technique can be applied in a number of ways. Preferably, the primary factor determining the way in which the technique can be used is the rate at which data in the log file is captured. If a web site is extremely popular, many thousands of requests may be captured every second and this obviously poses a problem for a process designed to analyse this information and produce a timely response. In this case, it is preferred to take periodic "snapshots" of the log file, whereby each snapshot contains a defined number of log entries. Then, the snapshots are processed and the results are added to a historical database recording the data.

Finally, if the output of the monitoring technique is that a session ID becomes associated (step 530) with a robot with a probability value, then processing passes to FIG 6 at "B" and the next stage of robot management is performed.

3. A robot management technique

The aim of this management technique is to proactively manage the presence of robots on a web site using pre-defined rules and responses based on robot behaviour patterns.

This technique will now be described with reference to FIG 6. A pre-requisite for this technique is the identification of one or more

robots via the activity logging technique and the data gathered about a robot's behaviour via the activity monitoring technique described above. The data associated with the identified robot and its behaviour is input into FIG 6 at "B".

5

Preferably, each robot detected on a system by the aforementioned monitoring technique is given (step 600) a dynamically generated "profile". The profile contains information about the behaviour and characteristics of a particular robot, and includes fields such as:

10

1. Current probability - that is, the value of "P" which is calculated as described above.

2. The time passed since the first detection of the robot

15

3. The number of unique hidden pages visited - a value is calculated which is a percentage of the total number of hidden pages

20

4. User agent - this is extracted from an HTTP request and can either be "consistent", that is, the same name is used for every request in which case the name is recorded, or "variable".

5. An IP address

25

Some of the fields can be constructed using data extracted from the log file, in this example: the number of unique hidden pages visited, user agent name and the IP address. The profile of a particular robot is used to determine how to respond to that robot's presence on a web site. A hierarchical series of "management rules" are defined to control this process and utilises (step 605) the fields. For example:

30

```
IF IP_ADDRESS IN BLACKLIST
THEN MANAGEMENT_RESPONSE = 1
```

35

```
ELSE
```

```
IF PAGES_VISITED <=10%
AND TIME_SINCE_DISCOVERY > 5 MINUTES
AND USER_AGENT = CONSISTENT
THEN MANAGEMENT_RESPONSE = 2
```

40

```
ELSE
```

```
IF PAGES_VISITED > 10%
```

```
AND TIME_SINCE_DISCOVERY <= 5 MINUTES
AND USER_AGENT = CONSISTENT
THEN MANAGEMENT_RESPONSE = 3
```

```
5 ELSE
```

```
MANAGEMENT_RESPONSE = 4
```

10 In the above example, the value defined for the "MANAGEMENT_RESPONSE" parameter indicates the response to be activated for managing a robot's presence on a web site.

15 In the example above, if the first rule is obeyed, that is, a positive result to step 610, the system activates (step 615) response "1". Rule 1 determines whether a robot is associated with a known blacklisted IP address. That is, an IP address of a potential robot is compared to a stored list of known IP addresses of blacklisted robots. If a match occurs, the response may simply be to deny any further requests from the robot, which can of course be identified via its session ID. If the first rule is not obeyed, that is, a negative result to step 610, processing passes to step 620.

25 In this step, the second rule deals with robots, which are slowly searching a web site. That is, the potential robot has searched less than or equal to ten percent of the web pages within a web site over a time period of more than five minutes. The potential robot in this example is additionally using a consistent user agent name. This would be typical of a robot, since on the whole, many individual users would use different user agent names.

30 If the second rule is obeyed, that is, a positive result to step 620, the system activates (step 625) response "2". An example of a response that would be initiated for these types of robots may be to do nothing other than to continue to monitor their behaviour. If the second rule is not obeyed, that is, a negative result to step 620, processing passes to step 630.

40 In this step, the third rule handles robots that are rapidly trawling through a web site. That is, the potential robot has searched more than ten percent of the web pages within a web site, over a time period of less than or equal to five minutes. The potential robot in this example is also using a consistent user agent name. These types of robot may be a potential hazard to a web site by consuming too many resources. If the third rule is

obeyed, that is, a positive result to step 630, the system activates (step 635) response "3".

One example of this response may be to withhold web pages from a requesting robot - if the session ID associated with that robot is stored, then any further requests made by that robot using the same session ID can be denied. If the third rule is not obeyed, that is, a negative result to step 630, processing passes to step 640.

In this step, the response "4" serves as a default selection or catch-all rule. The purpose of this rule is to deal with all robots not covered by the other rules. After each system response (steps 615, 625, 635, 640) processing finishes. Of course, FIG 6 is one example of a set of management rules. It should be understood that many other configurations of rules could be applied, as desired.

Advantageously, the management rules may be as simple or complex as required. Factors to be considered by an administrator when determining the rules, include the nature of a web site, the web site's susceptibility to robots or the perceived dangers or problems that the robots may provoke. Preferably, the logging, monitoring and management techniques are implemented on an application server.

In a further embodiment of the present invention, the management responses could be used to control the amount of server processing time allocated to individual robots. For example, a response may establish that:

"If a server is busy, reduce the priority of the thread controlling a robot's HTTP requests"

Alternatively, the response could establish that:

"If a server is not busy, then the priority of the thread controlling a robot's HTTP requests is raised"

Advantageously, the preferred embodiment ensures that a robot does not consume too much of a system's resources at the expense of other users. However, during slack times, such as when a server load is at its lowest, robots could be positively encouraged by minimising the response time to their HTTP requests. The preferred embodiment will therefore encourage "positive feedback" for intelligent systems. An example of an intelligent system is a robot that is able to detect when it is allowed to navigate a web site. An intelligent system could therefore identify the best time to

traverse a specific web site based on the response time from a server and will only navigate at optimum periods.

5 Preferably, a higher-level "global" set of management rules is used to control all robots using a web site. For example, these rules determine the number of threads that can be allocated to robots at any given time, the default priority allocated to robot threads and so forth.

10 It will be apparent from the above description that, by using the techniques of the preferred embodiment, an end-to-end process for automatically protecting a system from malicious robots is achieved, whilst also providing flexibility in the way in which the process can be implemented. The present invention is also advantageous in that the amount of server processing time allocated to individual robots can be controlled.

CLAIMS

1. A method of controlling access to information in a distributed data processing system having: a server for storing said information, in which
5 said server further comprises a logging tool for creating a log file, and a client computer comprising an application program for controlling a software agent, in which said software agent requests said information from said server, said method comprising the steps of:

10 identifying a software agent;

in response to said identifying step, storing all requests from said identified software agent in said log file;

15 in response to said storing step, analysing said log file;

in response to said analysing step, monitoring behaviour of said identified software agent, and

20 in response to said monitoring step, invoking at least one of a plurality of pre-defined rules to control said behaviour of said identified software agent.

2. A method as claimed in claim 1, in which said information is
25 represented within any number of a plurality of web pages, each of said any number of a plurality of web pages comprising a non-visible link.

3. A method as claimed in claim 2, in which a software agent requests
30 one of a plurality of web pages, said identifying step further comprises the steps of:

dynamically generating a first unique identifier;

35 dynamically inserting said first unique identifier into a non-visible link associated with said one of a plurality of web pages, and

determining whether said one of a plurality of web pages is associated with further of a plurality of web pages.

40 4. A method as claimed in claim 3, in which:

if said determining step is successful, said first identifier is dynamically inserted into further non-visible links.

5. A method as claimed in claim 3, in which if said determining step is not successful, said identifying step further comprises the steps of:

5 sending said one of a plurality of web pages to said identified software agent;

10 in response to said sending step, requesting, from said server by said identified software agent, any number of a plurality of links associated with said one of a plurality of web pages;

15 in response to said requesting step, extracting, by said identified software agent, said any number of a plurality of links;

20 in response to said extracting step, passing, by said identified software agent, said any number of a plurality of links to said client application program, and

25 in response to said passing step, determining, by said client application program, which of said any number of a plurality of links to display.

6. A method as claimed in claim 5, in which said any number of a plurality of links is displayed within a web browsing session running on said client computer.

7. A method as claimed in any preceding claim, in which said analysing step further comprises the step of:

30 identifying a first value associated with said any number of a plurality of web pages and a second value associated with said further of a plurality of web pages.

8. A method as claimed in claim 7, in which said monitoring step further comprises the steps of:

35 utilising said first and second values to generate a third value, in which said third value is associated with said identified software agent, and

40 utilising said third value and a fourth value associated with said all requests, to associate said first identifier and said identified software agent with a fifth value, in which said fifth value is associated with a probability.

9. A method as claimed in claim 8, in which if said third value is not more than or equal to a first pre-determined threshold, said log file is analysed further.

5 10. A method as claimed in claim 8, in which if said fourth value is not more than or equal to a second pre-determined threshold, said log file is analysed further.

10 11. A method as claimed in claim 7, in which said identified software agent is associated with a profile, said profile comprising any number of a plurality of data fields unique to said identified software agent, in which said invoking step further comprises the steps of:

15 associating said any number of a plurality of data fields with said at least one of a plurality of pre-defined rules;

determining whether a pre-defined response needs to be activated, and

20 in response to a successful determining step, activating said pre-defined response.

25 12. A method as claimed in claim 11, in which said log file further stores an address associated with a software agent and a name associated with said software agent.

13. A method as claimed in claim 11 or claim 12, in which at least one of said any number of a plurality of data fields is extracted from said log file.

30 14. A method as claimed in any preceding claim, in which a second unique identifier is generated and further inserted into a non-visible link.

35 15. A method as claimed in any preceding claim, in which said distributed data processing system further comprises an application server.

40 16. A method as claimed in any preceding claim, in which said at least one of a plurality of pre-defined rules controls a plurality of thread priorities associated with said server, in which at least one of a plurality of threads is associated with a software agent.

17. A system for controlling access to information, for use in a distributed data processing system, said distributed data processing system comprising: a server for storing said information, in which said server further comprises a logging tool for creating a log file, and a client

computer comprising an application program for controlling a software agent, in which said software agent requests said information from said server, said system for controlling access comprising:

5 means for identifying a software agent;

 means, responsive to said identifying means, for storing all requests from said identified software agent in said log file;

10 means, responsive to said storing means, for analysing said log file;

 means, responsive to said analysing means, for monitoring behaviour of said identified software agent, and

15 means, responsive to said monitoring means, for invoking at least one of a plurality of pre-defined rules to control said behaviour of said identified software agent.

18. A distributed data processing system comprising: a server for storing
20 said information, in which said server further comprises a logging tool for creating a log file, and a client computer comprising an application program for controlling a software agent, in which said software agent requests said information from said server, and

25 a system for controlling access to information, as claimed in claim 17.

19. A computer program product comprising computer program code stored on
30 a computer readable storage medium which, when executed on a computer, performs the method of any of claims 1 to 16.

ABSTRACT

A METHOD, SYSTEM AND COMPUTER PROGRAM FOR CONTROLLING ACCESS IN A
DISTRIBUTED DATA PROCESSING SYSTEM

Controlling access to information in a distributed data processing system. The distributed data processing system has a server, which stores the information and also further comprises a logging tool for creating a log file. The distributed data processing system also has a client computer comprising an application program for controlling a software agent. When the software agent requests information from the server, a process to identify the software agent is invoked. In response to the identification, all the requests from the identified software agent are stored in the log file and this data in the log file is analysed. The data is also utilised in the process of monitoring the behaviour of the identified software agent. In response to the monitoring process, at least one of a plurality of pre-defined rules is invoked, in order to control the behaviour of the identified software agent.

1/5




FIG. 1 PRIOR ART




FIG. 2 PRIOR ART


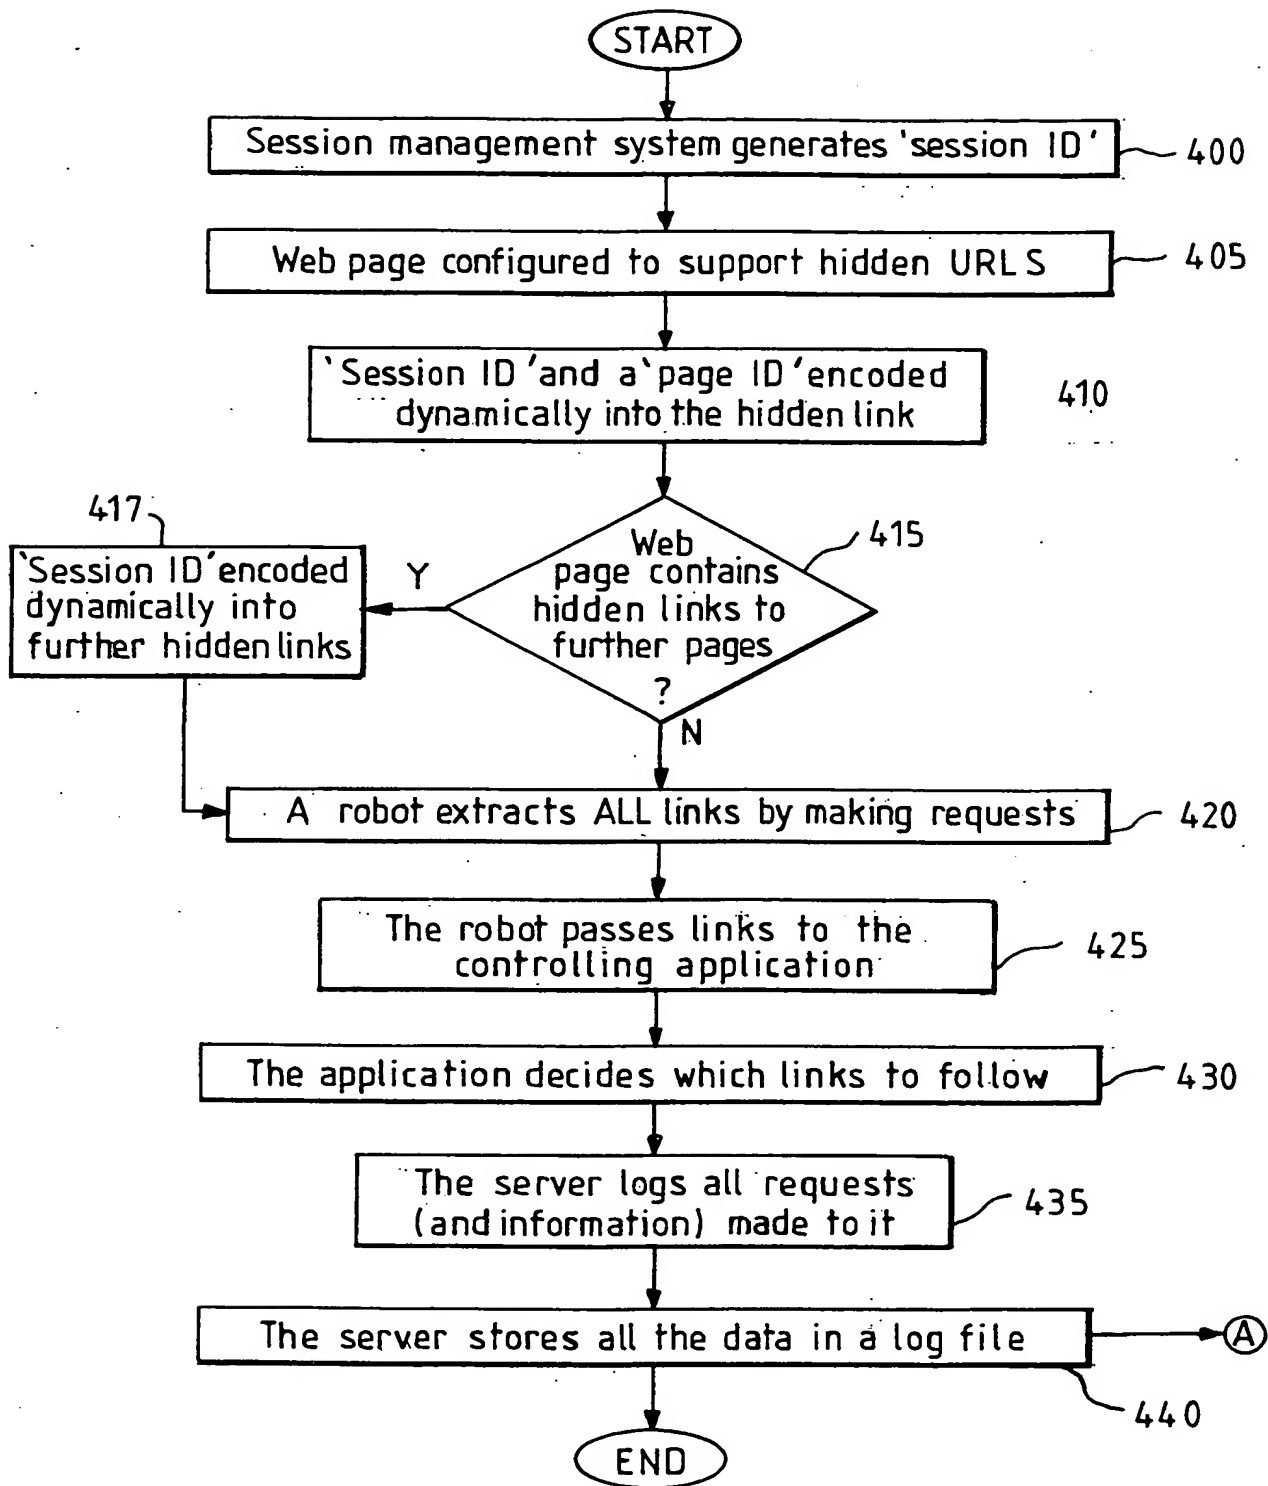
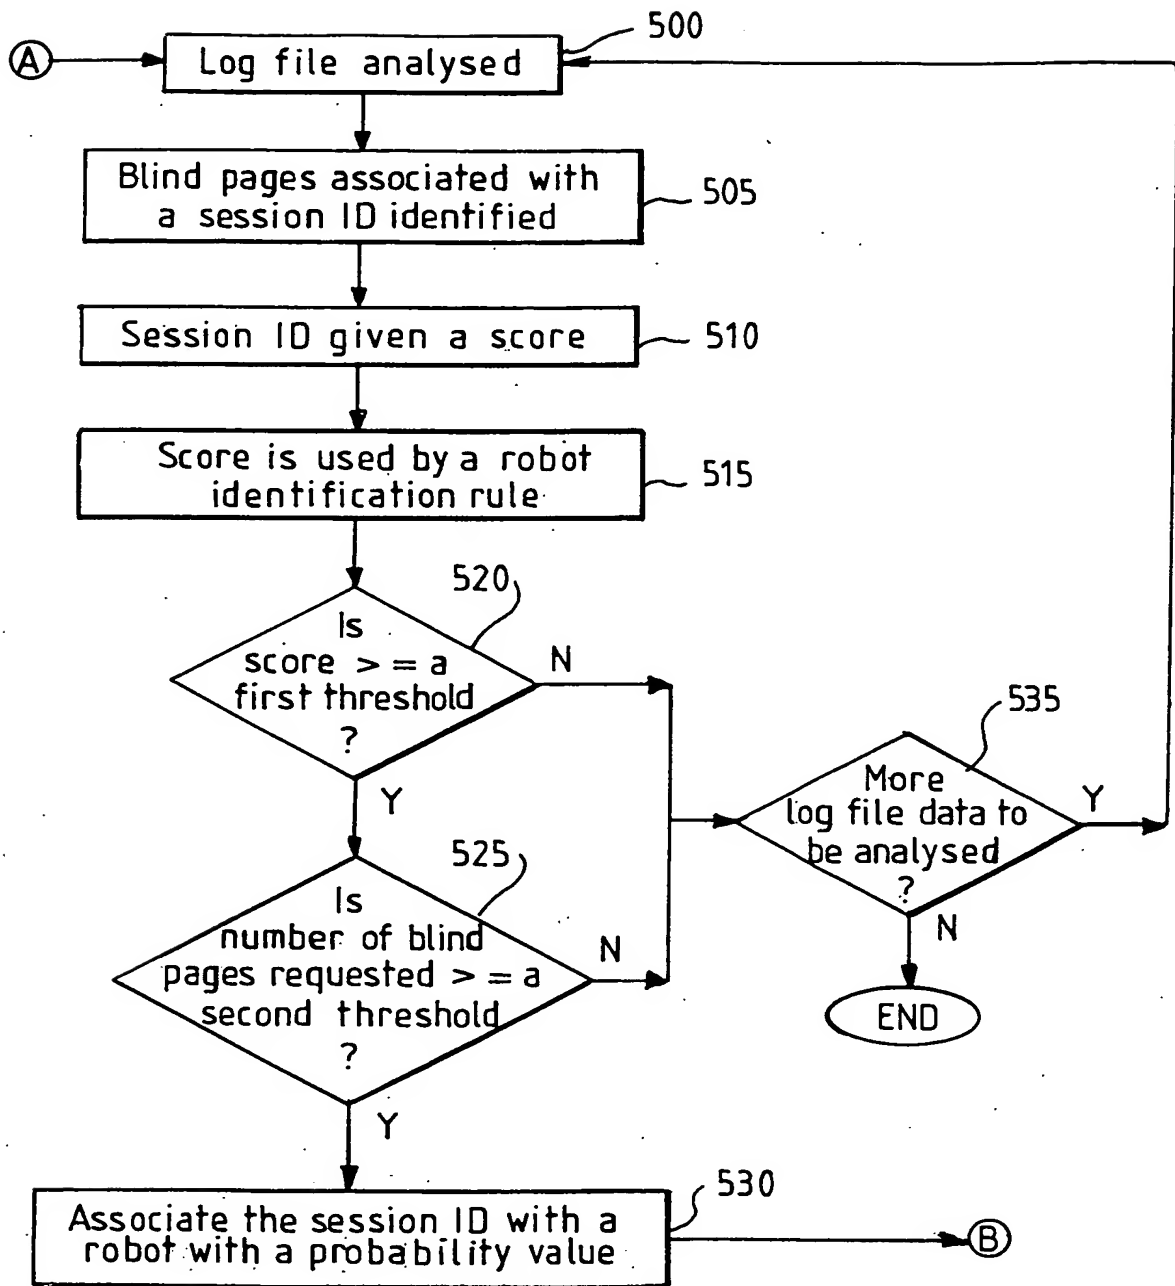


FIG. 3 PRIOR ART

3 / 5

FIG. 4

4 / 5

FIG. 5

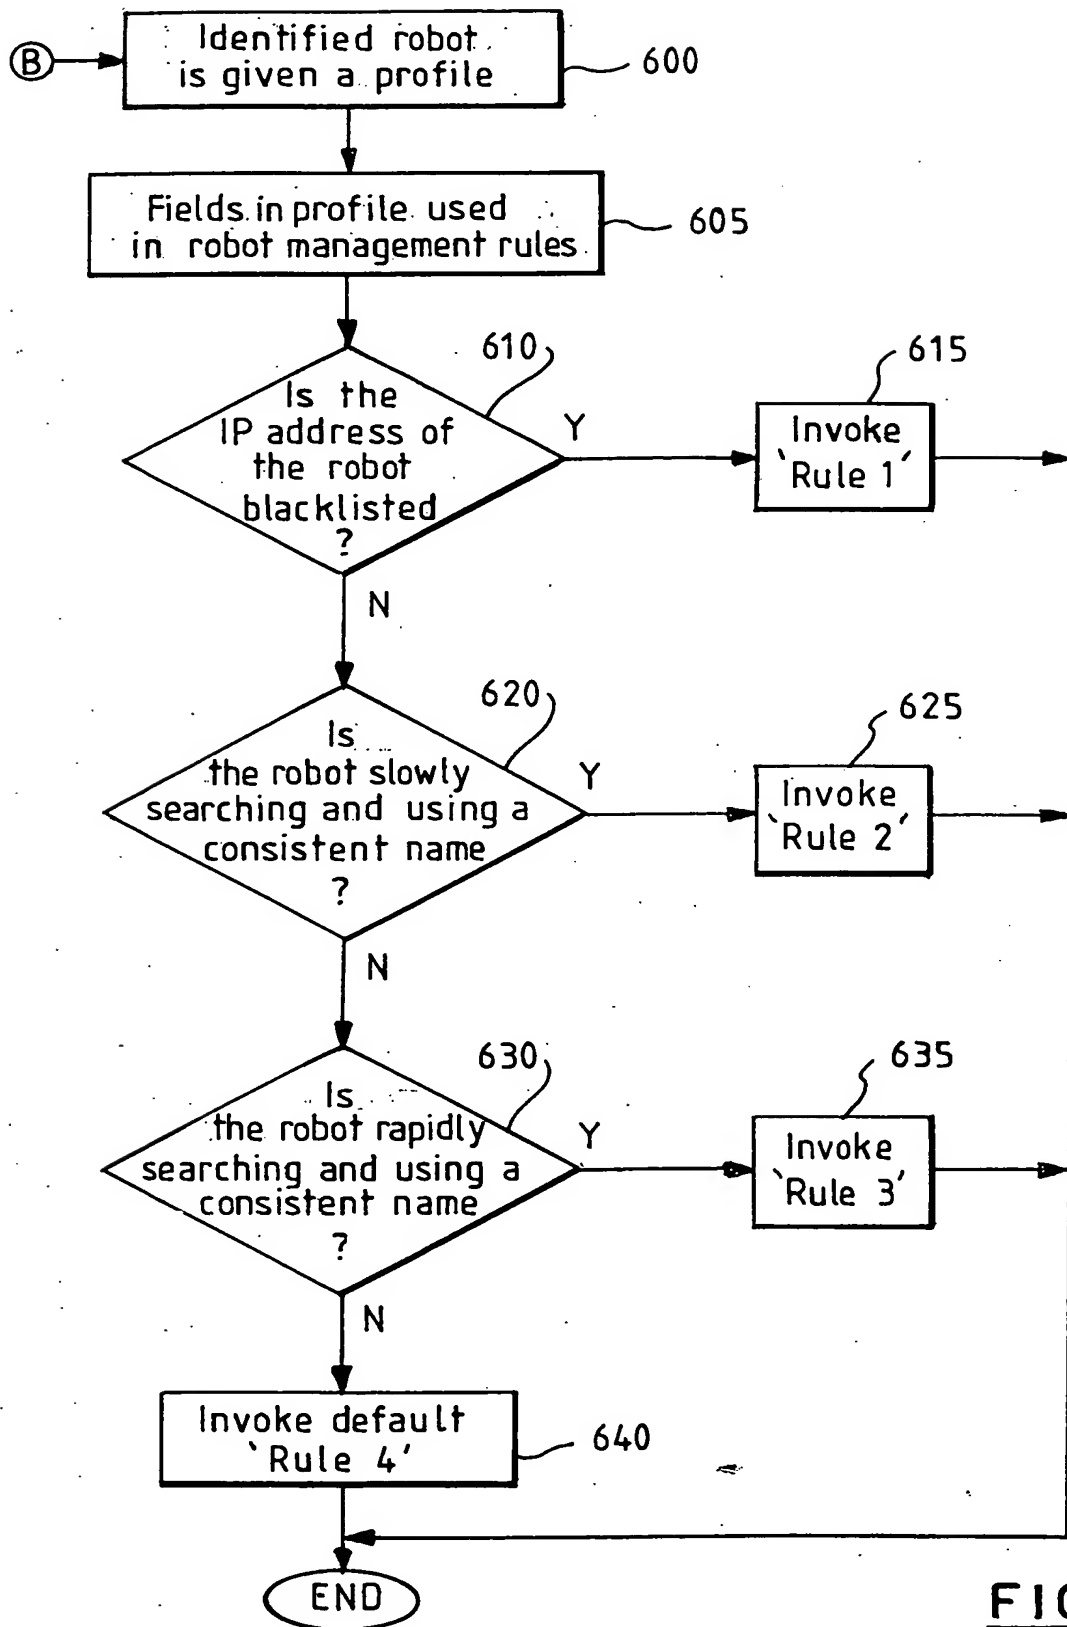


FIG. 6